Channel Access Synchronization Test

05/2001
P. Chu

Introduction

For most of the SNS applications, we need *time correlation* among many data from various controls and diagnosis systems. A straightforward way to correlate the data is to use the synchronous-get, *ca_sg_array_get*(), routine from the EPICS Channel Access Libraries. We can specify a small time-out window for synchronous data gathering, *i.e.* get the data we want within the time-out window. The purpose of this test is to learn if this method is suitable for our need.

How the test was done

The test was done via light network traffic (mostly off-hours). The program was written in C++ (source code is available on linux1:/work1/epics/) and the iteration was controlled by Shell Scripts which are available in both *bash* and *tclsh*. The tclsh version of the C++ program requires being compiled as a shared object and dynamically loaded into the shell during the run time. For comparison reason, I only list the results from the *bash* Script. The bash Script sent out 54,000 synchronous-get requests for time stamps of 2 PVs (also tested with 16 PVs) in 1 Hz rate for a total time of about 15 hours. In bash, 1 Hz is the fastest rate one can go with *sleep* command (in tclsh, one can go as fast as 1000 Hz with *after* command).

Results

The returned time stamps from the function calls were saved in files. After check the output files, I found the following 3 error categories:

- connection to any of the PVs failed within the time-out period,
- ca_sg_get() failed without the time-out period (note, this time-out window can be different from the previous one),
- "synchronous" PVs with different time stamps (if within the time-out window, it is acceptable).

The following is an example of the test result: This is for 2 PV time stamps with 10 msec time-out window at about 4% IOC CPU load.

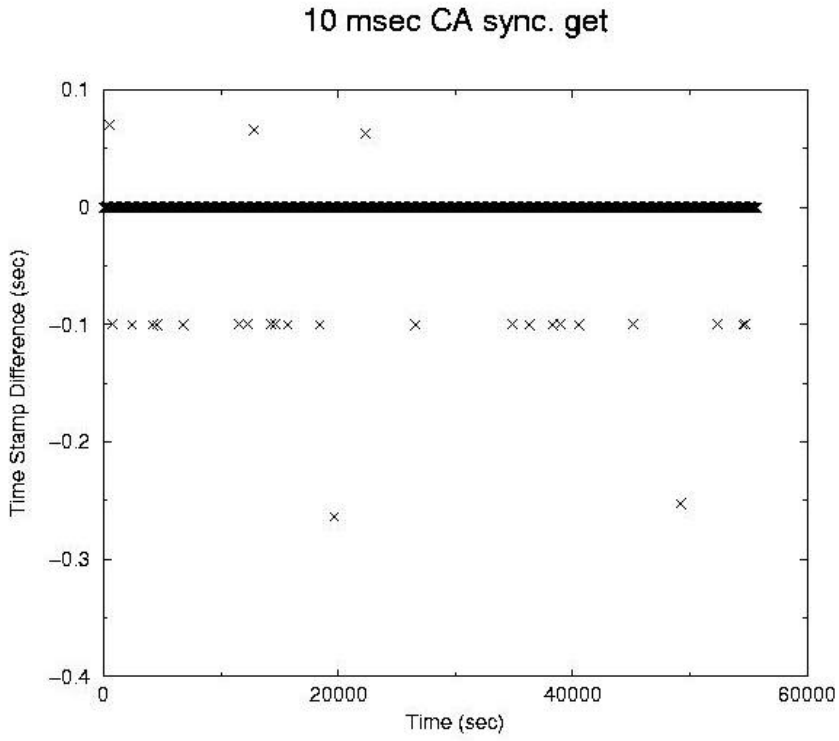

Fig.1.The horizontal axis is the total time (~54000 s) and the vertical axis is the time stamp difference (first – second)

The compiled result is plotted in Fig.2. In Fig.2, there are total 5 tests with various time-out windows and IOC CPU loading:

1: 10 ms window (CPU load ~4%)
2: 50 ms window (CPU load ~4%)
3: 100 ms window (CPU load ~4%)
4: 90 ms window (CPU load 60~80%)
5: 80 ms window (CPU load 60~80%)

Conclusion

As shown in Fig.2, the test results suggest that the IOC CPU load is the major factor that limits how fast the *ca_sg_array_get()* can go. When the

CPU load is higher, we have to enlarge the time-out window to keep a reasonable rate of successful synchronous-get. If the CPU load is reasonable (say, below 50%), the synchronous-get can go as fastest as 10 Hz with very few failures. Of course, the network and shell overhead should be taken into account for the real world practice.

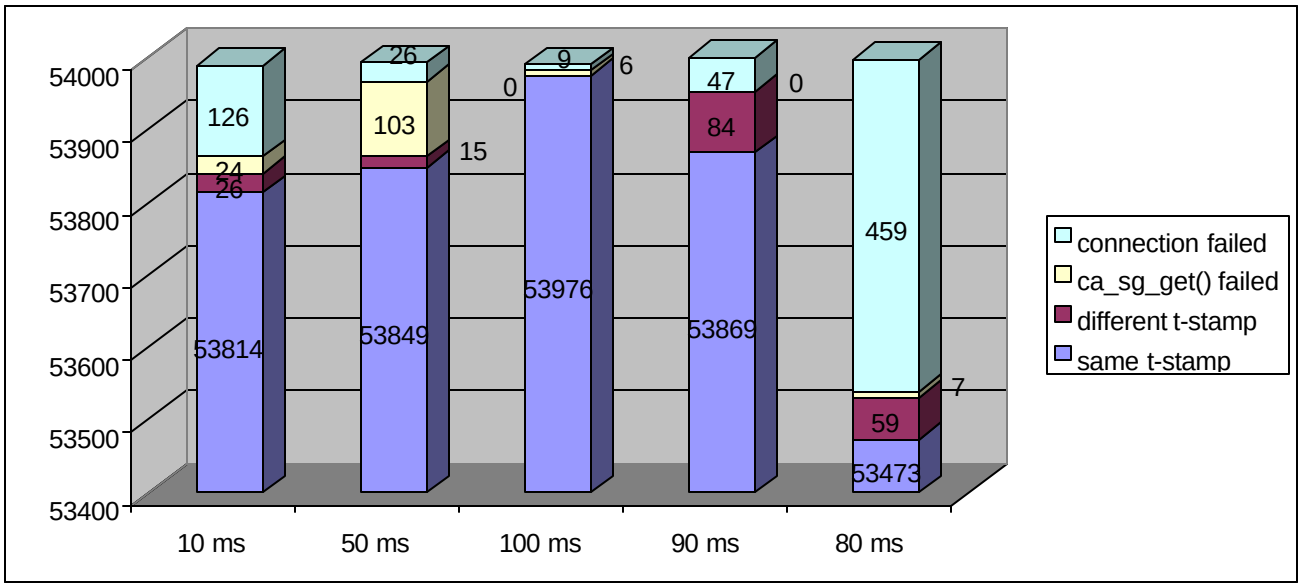

Fig.2. A compiled summary of 5 tests. The vertical axis is the number of various conditions.